AD-A084 930 ROYAL AIRCRAFT ESTABLISHMENT FARNBOROUGH (ENGLAND) F/G 8/2
GEOMETRIC CORRECTION OF SATELLITE IMAGERY.(U)
SEP 79 J M WILLIAMS
UNCLASSIFIED RAE-TR-79121 DRIC -BR-72781 NL

1 OF 1
AD
A084930

END
DATE
FILMED
7-80
DTIC

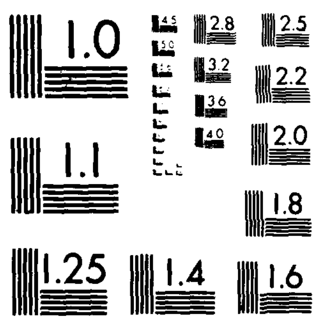

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

BR72701

UNLIMITED

TR 79121

LEVEL II

TR 79121

AD A084930

ROYAL AIRCRAFT ESTABLISHMENT

*

Technical Report 79121

September 1979

DTIC ELECTE APR 30 1980 S D E

GEOMETRIC CORRECTION OF SATELLITE IMAGERY

by

J.M. Williams

*

Procurement Executive, Ministry of Defence
Farnborough, Hants

UNLIMITED

UDC 535.317.1 : 771.537 : 629.195 : 339.5

ROYAL AIRCRAFT ESTABLISHMENT

(9) Technical Report 79121

Received for printing 17 September 1979 (11) 17 Sep 79

(6) GEOMETRIC CORRECTION OF SATELLITE IMAGERY

by

(10) J. M. Williams (12) 31

SUMMARY

Imagery from Landsat and other remote sensing satellites suffers geometric distortion which requires correction. This Report describes how ground control points can be used to determine the transformation between image coordinates and some known projection.

(14) RAE-TR-79121

(18) DRIC (19) BR-72782

Accession For	
NTIS GRA&I	✓
DDC TAB	
Unannounced	
Justification	
By	
Distribution/	
Availability Codes	
Dist. A	Avail and/or special

Departmental Reference: Space 569

310450

LIST OF CONTENTS

1 INTRODUCTION

This Report describes the geometric transformation of images received by remote sensing satellites such as Landsat[1]. These images are distorted due to satellite attitude and orbit variations, Earth curvature and rotation and sensor misalignments. The correction of these distortions would be relatively simple if they could be accurately modelled, but in practice the desired accuracy is not often achieved. Geometrically corrected images are needed for four purposes:

(1) to bring an image into a standard projection;

(2) to locate points of interest;

(3) to bring adjacent images into registration; and

(4) to overlay images of the same area from different sensors (though not necessarily in a standard projection).

Registration is a specialised case of transforming images into a standard projection, since if two adjacent images are transformed into the same projection they can then be joined together. The main aim was therefore to perform this transformation using ground control points (that is points in an image whose positions are accurately known in the required projection). These ground control points were located in the image and on a standard map, which for our purposes was a 1:50000 Ordnance Survey map[2]. The geometric transformation between the map and the image can then be used either to transform the complete image into the map projection, or to locate points in an image whose map coordinates are already known.

Investigations were carried out into suitable models for performing the transformations (see Appendix B), and into the various interpolation techniques which might be used to resample the transformed image (see Appendix D). Several programs were developed on the Prime 400 computer in the Remote Sensing Centre, and a broad outline of their structure and use is given in Appendix C. The problems of optimising computer usage are assessed in section 8 and the techniques for selecting ground control points are described in section 3. Methods of annotating images with crosses or other symbols and overlaying grids are discussed in section 6.

A brief description of the RAE image processing system is given in Appendix A. Finally, a set of photographic images are included. These are designed to show the effects of the various processing methods described.

2 LANDSAT IMAGES

The investigation was carried out using standard Landsat images, although any image in a similar format could have been used.

Landsat[1] is a NASA remote sensing satellite which is used to detect terrestrial radiation for use in image construction.

The images used originated from the Multi-Spectral Scanner (MSS), which records the intensity of reflected radiation in each of four spectral bands. The Earth's surface is broken up into overlapping 80 m squares and the net intensity of each square detected. The data collected is transmitted to ground stations for initial processing.

The images received by the RAE are recorded on Computer Compatible Tapes (CCTs) which have been pre-processed by the European Space Agency at the Earthnet station at Fucino (Italy). These images are stored in a standard format in the RAE processing system (see Appendix A). Each image consists of a set of picture elements (pixels), each of which corresponds to a rectangular area on the Earth's surface. The intensity level of each pixel is represented by an integer in the range 0 to 255 (that is 8 bits of information). This scale is nonlinear but it is at least possible to deduce the relative intensity of pixels.

An image can be thought of as having the pixels arranged in rows and columns as in its photographic representation, but an image itself is simply a collection of data. Anything manufactured from this data is not an image but an image product, in the same way that a photograph of an object is distinct from the object itself.

As each image may consist of up to 2500 rows with 3000 pixels in each row, a large amount of computer storage is needed. Since each pixel consists of an 8-bit word (byte) and an image contains four separate spectral bands, each complete image may require 30 megabytes (3×10^7 bytes) of storage capacity. During the following analysis only one spectral band was used , so references to an image should be taken as implying a single spectral band of one image. If desired, two or more bands could be identically corrected and used to form a composite image in colour. The enormous amount of data handling involved is one of the major problems associated with digital image processing.

3 SELECTION OF CONTROL POINTS

The method used to calculate the transformation needed to bring an image into a standard map projection involves locating a set of features in the image and on a map. To do this, the image must be displayed in a visible form so that

the location of features in pixel coordinates is possible, and it must then be decided which features are suitable for use as control points. A feature is accepted as suitable if it satisfies the following criteria:

(1) it is unambiguously identifiable in the image;

(2) the position in the image is known to the nearest pixel;

(3) the map reference is accurately known.

Many features in the image (*eg* fields, woods and open land) are very similar and are therefore unsuitable. Towns which show up well cover a vast number of pixels and, as specific parts of a town are usually difficult to identify, they are also unsuitable. Bodies of water, such as lakes, ponds, or reservoirs are easily recognisable, but are generally unsuitable because the level of water may vary considerably thus changing the outline. Some areas, such as gravel pits, may have been extensively altered since the last map revision. Major roads and rivers show up well on images and are often only one pixel wide. Thus the intersection of two roads, or a road and a river, may consist of only one pixel. The bulk of control points on an image will normally consist of this type of feature. The intersection of two motorways, however, should be viewed with caution, as junctions may be spread over several pixels. Occasionally main railway lines show up well enough for use, but as they often follow roads, it is difficult to find many intersections. Other suitable sources of control points are airfields, where the intersection of two runways can easily be found.

4 CALCULATING THE TRANSFORMATION

The transformation is determined by using a set of control points on the map and their counterparts in the image (see Fig 1a). This transformation is assumed to be a good approximation for the whole image. The geometric model used for the transformation is described in Appendix B.

The required image is displayed on the IDP3000 visual display system (see Appendix A) and suitable control points are selected. As many as possible are used, evenly spread over the image to give a representative sample. The digital cursor on the display is used to locate these points in pixel coordinates. The reading is taken twice and if the readings disagree the point is rejected as unsuitable. If only a few control points can be found each reading is made a few times and the mean value taken. The points are then located on the map; it is often convenient to take their map coordinates relative to a local origin near the centre of the image.

The program MATRIX (see Appendix C, section C.1) is then used to calculate the transformation matrix between the two sets of points.

The transformation used is between map coordinates and image coordinates, even though the image is to be brought into the map projection. The reason for this is that since the transformed image is constructed a line at a time, it is necessary to know where each of its pixels came from in the original image, in order to assign each the correct intensity level. Thus, each pixel of the image in the map projection must be mapped back to its original position (see Fig 1b). This is the inverse of the mapping shown in Fig 1a, which is from image coordinates into the map projection.

5 TRANSFORMING THE IMAGE

The transformation from map projection to image coordinates is now known. It would therefore seem a simple task to transform a complete image, but as previously noted, any type of image manipulation takes a large amount of processing time due to the amount of data handling required. Section 8 describes the program development carried out to lessen this problem.

The actual transformation is performed using the program TRANSFORMATION (see Appendix C, section C.2). This program uses the inverse mapping described above to transform the image into the desired projection. As previously mentioned, each pixel is mapped back to its original position. Generally this position will lie between four points in the original scene (see Fig 2a). The intensity assigned to the pixel should thus be a function of the intensity of these four points and their position relative to the transformed point. This raises the problem of finding a suitable interpolation formula. Appendix D describes the possible choices of such formulae and the ones most suited for used in this particular application.

The transformed image will not be exactly the same size as the original image, so that some of the pixels near the edges may not have been in the original scene (see Fig 2b). These pixels are left with an intensity of zero in the final image. To avoid this effect, a sub-image (typically 512 × 512 pixels) is often used. If this sub-image is suitably selected all its pixels map back into the original scene (see Fig 2c).

6 GRIDDING TECHNIQUES

As transforming a complete image can be a lengthy operation, alternative approaches to the problem were considered.

One method is to overlay the image with grid lines. These lines would be normal grid lines from a map, transformed into image coordinates. They will therefore not normally be straight but can still be used as a good approximation to a Cartesian system.

It is normal not to include a whole grid, but to mark grid intersections with crosses. These are transformed into image coordinates with the aid of a program TRANS.POINTS (see Appendix C, section C.4). The matrix for the transformation is obtained in exactly the same way as before (see section 4). Programs CROSS (see Appendix C, section C.5) and GRID (see Appendix C, section C.6) are then used to produce a grid-image. This is of the same dimensions as the original scene and consists of pixels with intensity 255 (background) and 0 (crosses). This can be used to produce a transparent photographic overlay on the Linoscan (see Appendix A). Photographic images contain registration marks in each corner, so the overlay can be located accurately on top of the original scene. The grid-image can also be merged with the original image, as explained below.

6.1 Annotation of images

Crosses can be added to the actual image, but this would obliterate those pixels in the image that lie under them, thus reducing the information content. The method used takes the pixels lying under a cross (which may be several pixels wide in order to be visible on a photograph) and modifies their value so that the cross is easily recognisable, but their intensity relative to adjacent pixels is preserved. This is achieved with the program OVERLAY (see Appendix C, section C.7). The half-level is defined to be the integer part of half the maximum possible pixel intensity (in this case 127). The original image and an image-grid are compared and those pixels coincident with a grid cross are modified as follows:

(1) if the pixel value is less than, or equal to, the half-level, the pixel value is incremented by the half-level, and

(2) if the pixel value is greater than the half-level, the half-level is subtracted from it.

That is, the pixel is replaced by its diametric opposite (mod 255). Thus crosses will be at a significantly different intensity level from nearby pixels, but intensity variations within the cross will permit the information content to be preserved. Fig 3b shows a typical sub-image annotated with crosses. The original sub-image can be seen in Fig 3a.

7 LOCATING POINTS OF INTEREST

As previously mentioned, certain points on a Landsat image are readily identifiable, whilst for others positive identification may be difficult. Methods involving geometric correction can be used to locate such points. Firstly the desired points are located on the map and their reference coordinates noted. The standard procedure is used to determine the transformation from the map projection into the image. The program TRANS.POINTS, used previously to transform grid points, can then be used to transform these points into pixel coordinates. If the image is then displayed on the IDP3000, the digital cursor can be used to locate the points.

If desired, the process used to annotate an image with crosses can also be used. Thus, an annotated image, or image overlay can be produced, identifying points of interest with crosses or other symbols (it is probably best to use small circles, since the points are not spaced regularly). This can be used to produce an image with the control points used identified on it.

8 COMPUTING PROBLEMS

It has already been mentioned that the large amount of data handling needed in image processing can use up a great deal of computer time. Computer time falls into two categories, namely:

(1) central processor time (CPU-time), *ie* the time spent in actual computation; and

(2) disc input-output time (disc I-O time), *ie* the time spent fetching data to and from disc.

In image work, the second of these is often the most important since images are stored on disc in direct access memory (DAM) files. Complete lines of data can be read from these files, but the time spent fetching the data can still be large compared with the processing time. If a line of data had to be read off disc for every pixel in the output image, the time used would be immense. This process would also be very wasteful since two adjacent pixels are often transformed into the same line, making it unnecessary to read the same line off disc twice in a row. A cache memory system to optimise disc reading is outlined in section 8.1.

Central processor time has been minimised by selecting the two simplest interpolation techniques, nearest neighbour and linear interpolation (see Appendix D).

Tests were run on a subscene (512 × 512 pixels) of a complete image. This subscene was rotated through 0, 1, 2, 5 and 10 degrees about its centre (0 degrees being the identity transform). Rotations were chosen because they are easily represented by a single parameter (as long as a consistent centre is used). Each rotation was carried out using both interpolation methods and the following statistics were recorded: central processor time used, disc I-O time used, and miss ratio (a parameter which represents how often a pixel value had to be read off disc).

8.1 The cache memory system

It is possible to store about ten lines of the image in local memory. This is only a small proportion of the complete image and so if this memory is to be used effectively, some method of selectively updating lines is necessary. In the cache memory system a record is kept of the numbers of the lines in memory and the order in which they were read in. When a pixel is transformed, its line number is checked to see if the line is in memory. If it is, it can be read directly from memory without any disc access. This is a considerably faster method of data acquisition. If the line is not in memory it is read off the disc and used to replace that line in memory which was read earliest. This process is very useful for small translations and rotations, since the same line may be required several times in a row. This system can also be used more efficiently if alternate lines are scanned in opposite directions when carrying out the transformation. The effect of this scanning technique is shown in Figs 4a and 4b, which show the path followed in the original image. If scanning is done as in Fig 4c the path followed in the original image will contain large jumps at the end of each line to a line not in memory (see Fig 4d). Every time a required line is in memory it is deemed a hit, a line which has to be read off disc is a miss. The ratio of misses to total lines is the miss ratio referrred to below. If the process did not use a cache memory, the miss ratio would be 100%.

8.2 Results

The table below shows the results obtained (all times are given in seconds).

Linear interpolation				Nearest neighbour			
Rotation	CPU	Disc	Total	CPU	Disc	Total	Miss-ratio (%)
0	375	23	398	143	25	168	0.05
1	443	194	637	192	138	330	0.39
2	535	410	945	290	489	779	0.83
5	801	989	1790	539	897	1436	2.13
10	1257	2154	3411	984	2064	3048	4.28

Fig 5 shows the total time and CPU-time for both methods, showing how with larger transformations the disc time begins to dominate the total time used, until with large rotations, nearest neighbour is not significantly faster. Fig 6 shows how the miss ratio increases with larger rotations, demonstrating that even with sizeable rotations, the cache memory system vastly reduces disc I-O time. Since few of the transformations will involve rotations larger than 10 degrees, this process is very effective, because it reduces the disc I-O time to the same level as CPU time.

9 CONCLUSIONS

The conclusions of this study can be summarised as follows:

(1) Although the geometric transformation of digital imagery is very time-consuming, it is possible to produce an effective system without specialised hardware.

(2) The correct selection of a large number of control points is an essential pre-requisite for determining accurate transformations.

(3) For less accurate work, the superimposition of grids on images can prove very useful, especially for locating points of interest.

(4) Different interpolation processes have different effects on both the transformation time and the image quality. In the present system only the simpler methods were practicable.

(5) The processing could be speeded up by the use of an array processor or specialised interpolation hardware.

(6) Techniques for the automatic location of control points (by statistical correlation with typical features) could prove useful for large-scale production work: these should be the subject of future investigations.

Appendix A

THE RAE PROCESSING SYSTEM

The remote sensing section has developed a system for processing images from various sources into a standard format. These images are stored on magnetic tape and are of various types, *eg* Multi-Spectral Scanner (MSS), Return Beam Videcon (RBV), Synthetic Aperture Radar (SAR) etc; for each data source a program has been developed which reads the tapes into a Direct Access Memory (DAM) file. These files are written in the same format for all data sources. They consist of 80 characters of text, followed by the image parameters, number of pixels per line and number of lines per image. Finally the actual data (pixel intensity values) are stored as either 8 or 16 bit words.

A suite of software has been developed to perform various processing jobs on the image files. These include contrast stretching, edge enhancement and principal components analysis.

A Linoscan image writer can be used to produce photographic representations of images (see Fig 3). This machine is connected to a Prime-200 computer. The required images are written to magnetic tape in a standard format. This tape is then played back to the Prime-200 and used by the Linoscan to produce a monochromatic negative.

The system also has a Plessey IDP3000 visual display unit coupled to a Prime-300 computer. Input to this system is also from magnetic tape. The display can be used to produce false colour composites of images by overlaying linear combinations of different spectral bands and assigning each one to a different colour. It also has a digital cursor which can be used to determine the coordinates of a particular pixel.

Appendix B

SELECTION OF A SUITABLE TRANSFORMATION

An image can be mapped into a given projection by a one-to-one transformation. If the image were produced on a 'rubber sheet' it could be continuously stretched and deformed until it fitted the projection. The transformation function can be approximated by a polynomial in two variables so that if a point has coordinates (x_1, x_2) in one system and (y_1, y_2) in a second, then

$$y_1 = \sum_{i=0}^{n} \sum_{j=0}^{n} a_{ij} (x_1)^i (x_2)^j$$

and

$$y_2 = \sum_{i=0}^{n} \sum_{j=0}^{n} b_{ij} (x_1)^i (x_2)^j$$

where a and b are constants and n is the order of the equation. This polynomial is a truncated version of the local two-dimensional Taylor polynomial (expanded about the centre of the transformation). As the order of this polynomial is increased it becomes a closer approximation to the actual transformation function. In the present analysis a linear function was used, this provides a good first approximation to the transformation function and can be quickly evaluated on the computer. It is represented by the equation:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \end{pmatrix} \begin{pmatrix} 1 \\ x_1 \\ x_2 \end{pmatrix} .$$

The determinant $\begin{vmatrix} m_{12} & m_{13} \\ m_{22} & m_{23} \end{vmatrix}$ represents the increase in area under this transformation, since m_{11} and m_{21} only serve to translate the image without distorting its shape. In order to produce a transformed image which is the same size as the original one should divide the matrix by the positive square root of this determinant. This produces an image which is in correct map coordinates but with a different scale.

If one has several pairs of control points in the two projections, a set of equations in the six unknown matrix coefficients can be set up. If the number of pairs of points is greater than three the equations are over determined and have no exact solutions.

The transformation can be represented by the equation $\underline{y} = M\underline{x}$ where $\underline{x}$ and $\underline{y}$ are linear vectors and M is the matrix to be determined. This equation can be solved by least squares methods. This involves finding the coefficients which minimise the function

$$S = \sum_{i=1}^{n} (\underline{y}_i - M\underline{x}_i)^2 .$$

Appendix C

SOFTWARE DEVELOPED

A series of programs has been developed for geometrically transforming imagery. These programs were written in FORTRAN and are designed for use within the image processing system on the Prime-400 computer. Fig 7 shows how these various programs link together. The circles represent files or other storage media and the programs should be regarded as methods of accessing, altering and storing data. The programs could be adapted for use on other machines by merely altering the file-handling subroutines. The general structure of the programs is outlined below.

C.1 MATRIX

The image and corresponding map coordinates of the ground control points are stored in sequential access memory (SAM) files on the computer. These coordinates are used to set up the linear equations in the matrix coefficients, and these equations are then solved by the least squares method. The matrix is written to a SAM file for use in subsequent programs.

C.2 TRANSFORMATION

This program, which performs the actual transformations, utilises a cache memory system in order to speed disc access. The transformation matrix is read from a SAM file and applied to the original image. The transformed image is created sequentially as described in section 8.1. Each point in the output image is mapped back to its original position, and the intensity values of the four surrounding pixels are used to perform two-dimensional linear interpolation. These pixel intensity values are read from the image-file if they are not in local memory. When a completely corrected output line has been formed, it is written to the output image-file.

C.3 NEIGHBOUR

This program is the same as TRANSFORMATION, except that it uses nearest neighbour interpolation. Both programs could easily be adapted for different interpolation techniques or higher order transformation polynominals.

C.4 TRANS.POINTS

This program reads the transformation matrix and applies it to a set of points stored in a SAM-file.

C.5 CROSS

This program is used to produce a file of symbols. The symbole required is stored in a SAM file which contains the pixel positions that the symbol would have if centred at the origin. The centre points to be used are also stored in a SAM file. A set of points is made by adding all the pixel positions in the symbol to all the centre-points in turn. These new points are written to a SAM file and represent the positions of all the pixels located under the symbol. This output file must be sorted into ascending order with the Prime utility SORT.

C.6 GRID

This program is used to produce a grid overlay of a specified size. The file of sorted points produced by CROSS is read and an image file constructed. Every point on the grid becomes a pixel of intensity 0, whereas all other points become pixels of intensity 255.

C.7 OVERLAY

This program combines a grid with an image. The grid line and the image line are read simultaneously and the value of each pixel in the output line is determined using the half-level criteria described in section 6.1. This new line is then written to an image-file and the process repeated until the annotated image is complete.

Appendix D

INTERPOLATION TECHNIQUES

It is often desired to reconstruct a continuous function, given its value at certain discrete sample point. The process of estimating the function values at intermediate points is known as interpolation. In this process it is usual to assume that sample points are spaced at unit intervals (or in two dimensions at the vertices of unit squares). The interpolation functions considered are generally only used in one dimension. All of them, however, can be adapted for two-dimensional use as follows. The function is first used to interpolate along a line to produce a new set of sample points, and these sample points are then used to interpolate between lines.

Four standard interpolation techniques and one interesting new one, were considered for possible use. The criteria used to determine whether a particular method was suitable were:

(1) the cost in terms of computer time needed, and

(2) the preservation of the spectral properties of the image (it should be noted that if resolution is diminished the image will appear blurred).

The method of analysing computer time has already been discussed in section 8, so only a broad description of the relative speeds is given. The preservation of spectral quality was determined by calculating the Fourier transform of the interpolating function. This transform gives the distribution of energy against frequency for a given interpolation process. If this distribution has a small amplitude for higher frequencies, the image will appear blurred. For further background information and an explanation of theorems stated without proof, see Ref 3. The ideal Fourier transform would be a block of constant value in the range [-1, 1], see Fig 9b. Sampling theory states that the interpolating function with such a transform can be used to exactly reconstruct a band-limited (*ie* finite range) function from sample points. The Fourier transform of a function $f(x)$ is given by:

$$\hat{C}(\omega) = \int_{-\infty}^{\omega} f(x) \exp(-i\omega x)dx$$

where $i = \sqrt{-1}$ and $\omega = (\text{frequency}/2\pi)$.

If the original interpolating function is band-limited, the Fourier transform is of infinite extent. So any interpolation using a finite number of sample points will always produce attenuation of the higher frequencies and hence some degradation of the final image.

D.1 Nearest NEIGHBOUR

In this method the transformed pixel is assigned the intensity value of its nearest NEIGHBOUR in the original scene. As this only involves integer rounding it requires little computer time. This interpolation is equivalent to convolving with a block of constant height in the range $[-1/2, 1/2]$, see Fig 8a. The idea of convolving can be represented graphically as follows: the origin of the interpolating function is slid down the real number line until it coincides with the point at which the sampled function is to be estimated. The estimated value is found by multiplying the values of the interpolating function at the same points, with the value of the sampled function and summing these products for all sample points. Thus with nearest NEIGHBOUR the block (or kernel) can only lie over one sample point giving the estimated value of the sampled function, as that of the value at the nearest sampled point multiplied by one. The Fourier transform of this function is $(\sin \pi x)/\pi x$, see Fig 8b.

D.2 Linear interpolation

Linear interpolation is equivalent to drawing a straight line between two adjacent sample points (*ie* it assumes the function can be approximated by a series of linear functions). If the function $f(x)$ is sampled at $x = a$ and $x = a + 1$, then its value at $x = a + \theta$ $(0 < \theta < 1)$ is given by

$$f(a + \theta) = f(a) + \theta[f(a + 1) - f(a)] \ .$$

This method is also economic in computer time, since it involves only a few simple arithmetic operations. It is, however, slower than the nearest NEIGHBOUR method.

Linear interpolation is equivalent to convolving a function with a triangle whose vertices are (-1, 0), (0, 1) and (1, 0), see Fig 8c. The Fourier transform of this triangle is $[(\sin \pi x)/\pi x]^2 x$, see Fig 8d. Since this is the square of a function whose maximum amplitude is one, its value is smaller than that function for nearest NEIGHBOUR. It therefore attenuates higher frequencies more than nearest NEIGHBOUR interpolation and so significantly degrades the resolution of the interpolated image.

D.3 Ideal interpolation

The Fourier transform of $(\sin \pi x)/\pi x$ is a block, see Fig 9a&b. This function could be used to reconstruct the sampled function, but with a finite number of sample points it must be truncated and assumed to be zero outside a certain range. This appreciably alters the Fourier transform from a block. Shlien has produced several graphs demonstrating the effect of truncation outside various limits. Typical of these is the eight point truncation reproduced in Fig 9c. It should be noted that at the truncation points the derivative is non-zero, so that the function is not smoothly continuous. This discontinuity in the functional derivative leads to a distinctly jagged Fourier transform, see Fig 9d, and this in turn will lead to image degradation due to uneven attenuation, even in the lower frequencies. Even a severly truncated version is very expensive in computer time. The next two methods use this function as a basis for improved interpolation techniques.

D.4 Cubic convolution

Cubic convolution involves approximating $(\sin \pi x/\pi x)$ by a series of polynomials between sample points. To make two polynomials smoothly continuous at the sample points both the function values and the functional derivatives must be the same. The lowest order polynomial which can satisfy these objectives is a cubic (if higher order polynomials are used, it is normal to equate the first few derivatives at the sample points in order to uniquely define the polynomial). The approximating function, $f(x)$, must also have the same zeros and turning points as $(\sin \pi x)/\pi x$. If a cubic is fitted to the five sample points (-2, -1, 0, 1, 2), it yields the following function

$$f(x) = x^3 - 2x^2 + 1 \qquad 0 \leqslant x \leqslant 1 ,$$

$$f(x) = x^3 + 5x^2 - 8x + 4 \qquad 1 \leqslant x \leqslant 2 ,$$

$$f(x) = 0 \qquad |x| > 2 ,$$

and

$$f(-x) = f(x) .$$

Although this function is a very good approximation to $(\sin \pi x)/\pi x$, it suffers from the disadvantage of not having zero derivatives at the end of its range, see Fig 10a. The transform, see Fig 10b, shows that it amplifies higher frequencies, but not nearly as much as either nearest NEIGHBOUR or linear interpolation. It is however more expensive in computer time.

D.5 Shlien interpolation

Shlien interpolation is a method of adapting the function so that its derivative is zero at the extremes of the range. The polynomial which is unity at $x = 0$ and passes through 0 at $x = 1, 2, 3, \ldots n$ is given by

$$f(x) = \left(1 - \frac{x^2}{1^2}\right)\left(1 - \frac{x^2}{2^2}\right) \ldots \left(1 - \frac{x^2}{h^2}\right).$$

This is equivalent to expanding $(\sin \pi x)/\pi x$ as an infinite product and truncating at n terms, see Ref 3. Starting from this basis Shlien developed a method of producing a smoothed interpolation function, namely:

$$f(x) = \left[\frac{\sin \pi x}{\pi x}\right]\left[1 - \frac{x^2}{h^2}\right] . \qquad (h = 8)$$

Both this function and its derivative are zero at $x = \pm n$. The function and its Fourier transform are shown in Fig 10c&d. The Fourier transform is an excellent approximation to the block function, producing negligible attenuation of higher frequencies. This is the best interpolation technique of all those considered, albeit by far the most expensive in computer time.

It was considered that only nearest NEIGHBOUR or linear interpolation (bilinear for a two dimensional case) were practicable with the computer time available. Fig 3c&d show the effect of these two methods on an actual image which has been rotated through 20 degrees. To appreciate the difference it is necessary to examine the negatives under a microscope. As linear interpolation is an averaging method it tends to round off the higher and lower extremes of pixel intensity, thus producing a blurred picture. Nearest NEIGHBOUR, although not as accurate does at least produce a sharper picture. It would have been worthwhile to produce a Shlien interpolated image but the available hardware did not make this a viable proposition.

REFERENCES

No.	Author	
1	NASA	Landsat Data Users Handbook. Goddard Space Flight Center, Document No.76DS4258 (1976)
2	Ordnance Survey	Some facts and figures relating to the transverse mercator projection and the national grid. Leaflet No.37 (1972)
3	R.W. Hamming	*Numerical methods for scientists and engineers.* McGraw-Hill (1973)
4	Seymour Shlien	Geometric correction, registration and resampling of Landsat imagery. Canada Institute for Remote Sensing, Ottawa, Canada

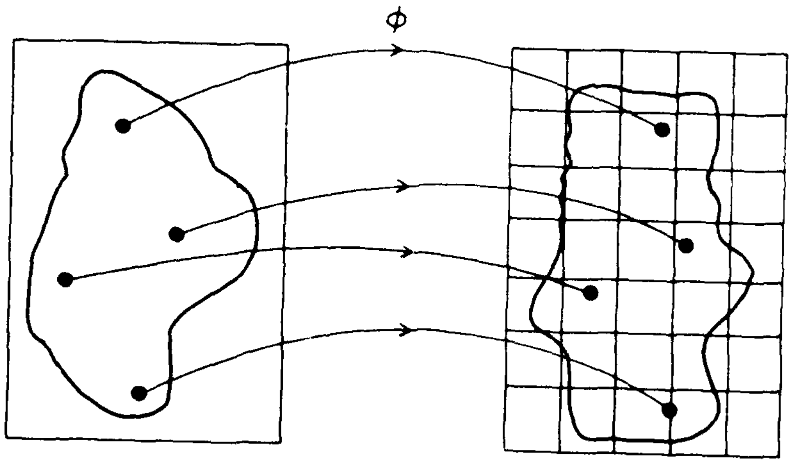

Fig 1a Control point mapping

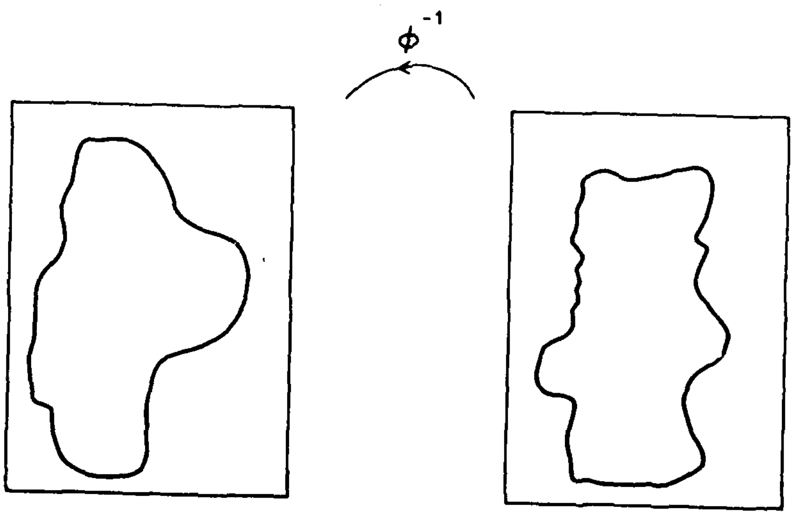

Fig 1b The inverse mapping

Fig 2a-c

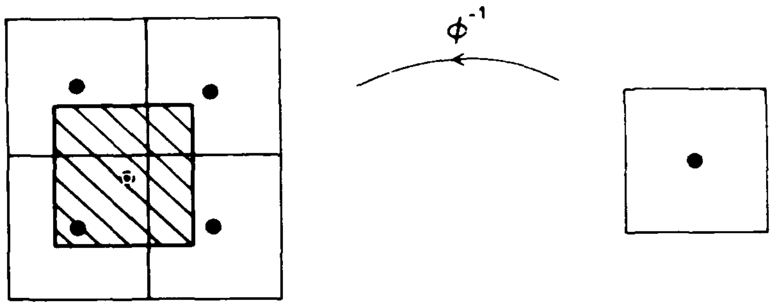

Fig 2a   The interpolation problem

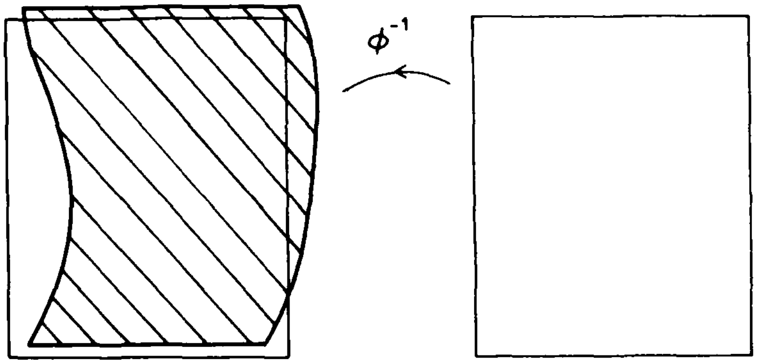

Fig 2b   The size problem

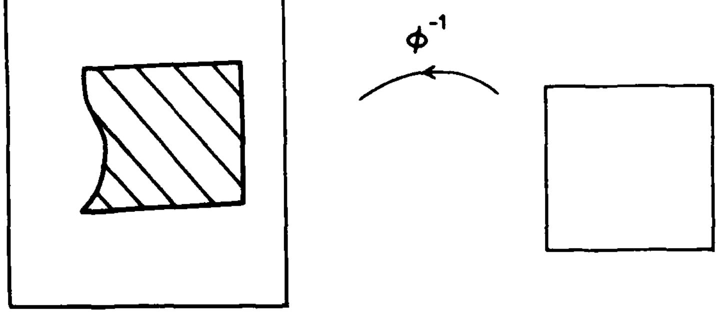

Fig 2c   Transforming a sub-image

TR 79121

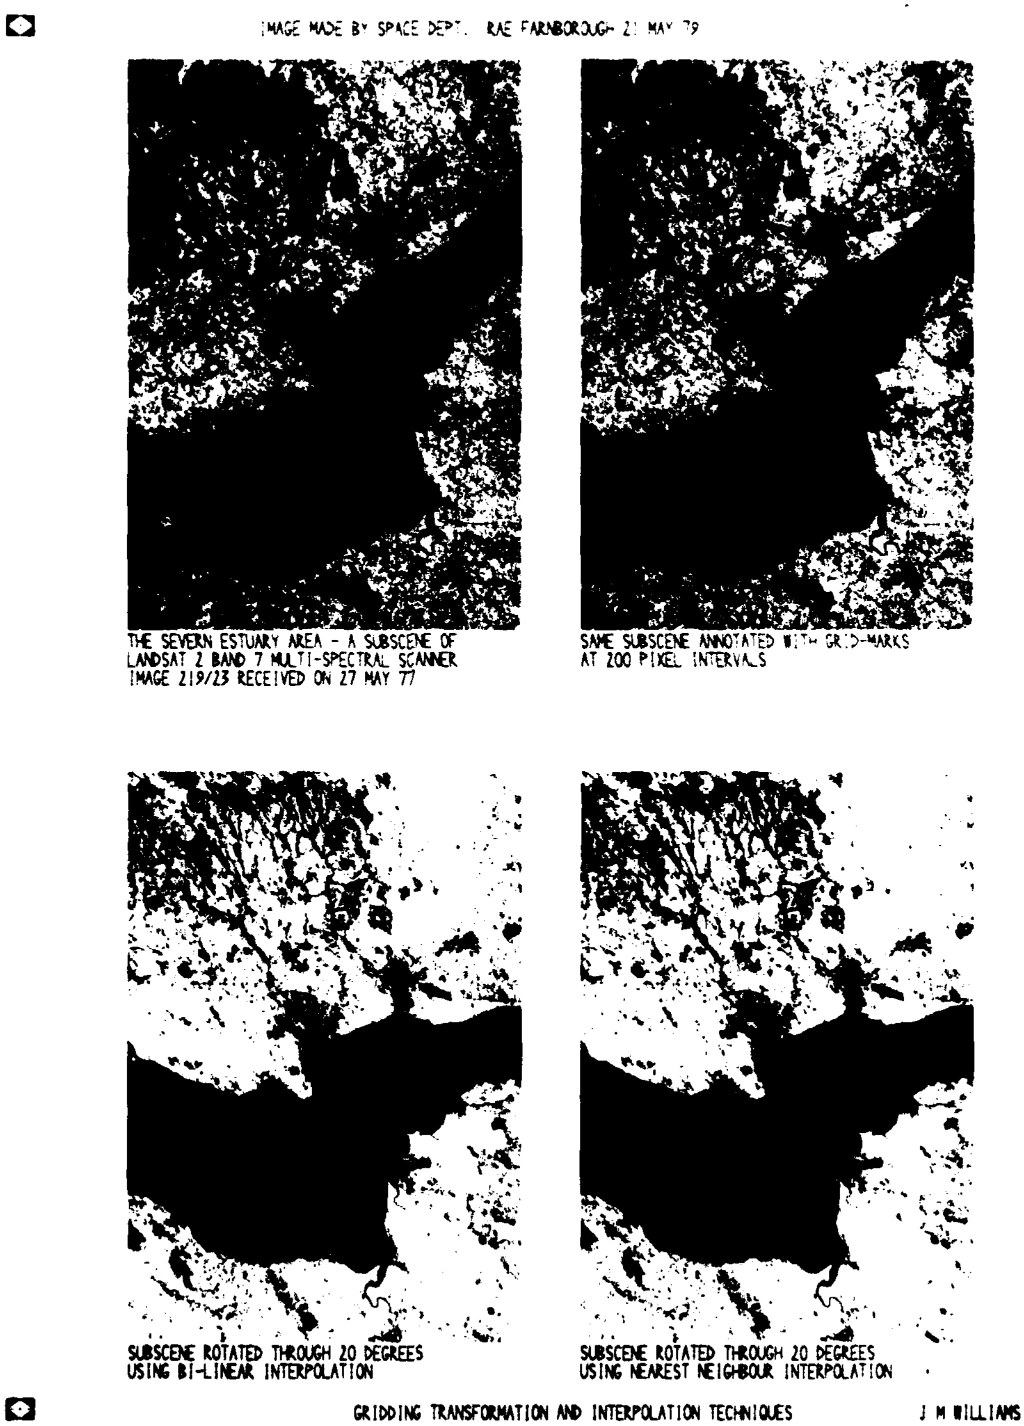

Fig 3 Gridding transformation and interpolation techniques

Fig 4a-d

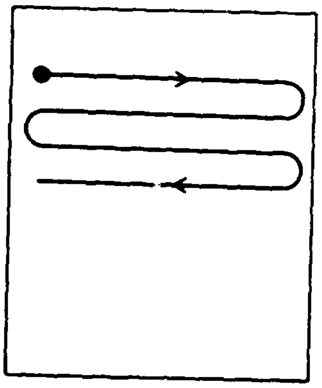

Fig 4a Alternate scan

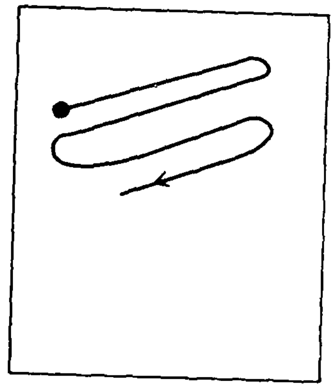

Fig 4b Path followed

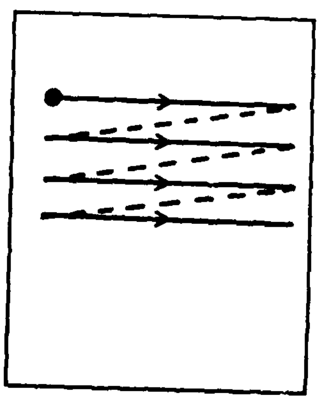

Fig 4c Linear scan

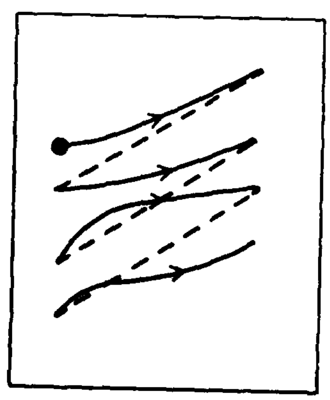

Fig 4d Path followed

TR 79121

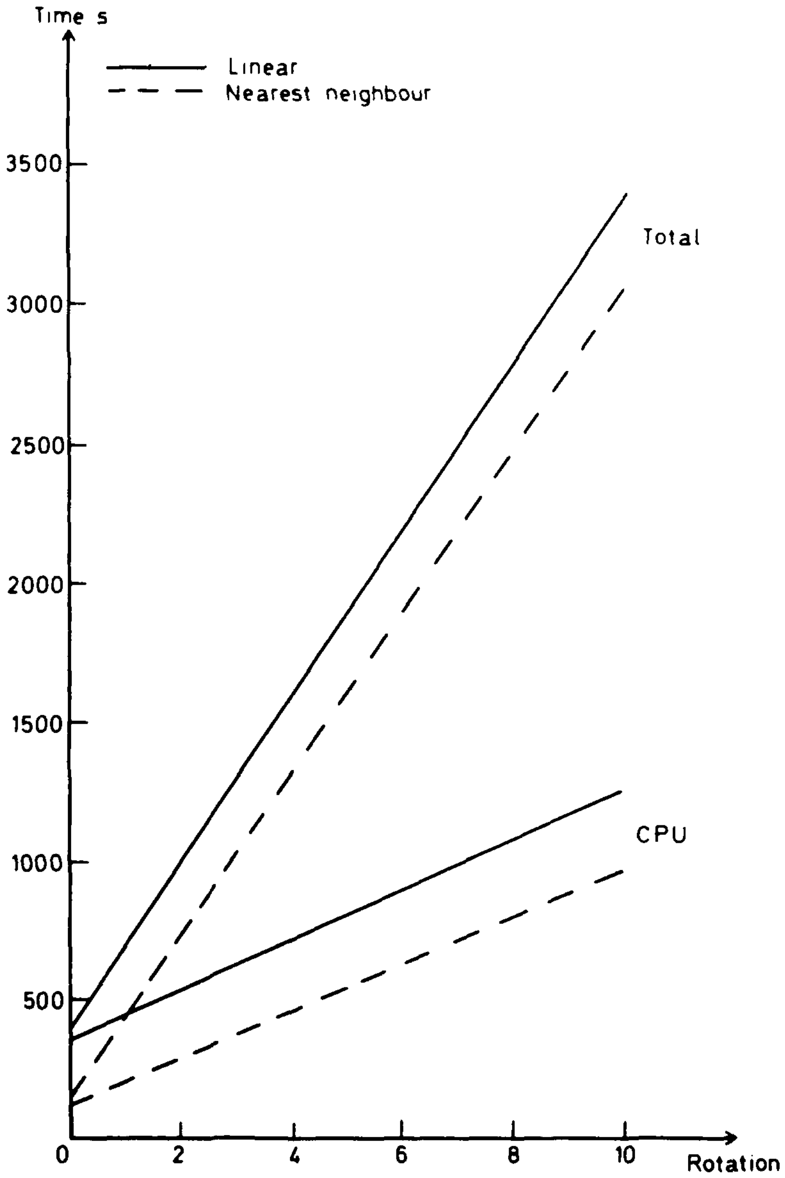

TR 79121

Fig 5 Transformation timings

Fig 6

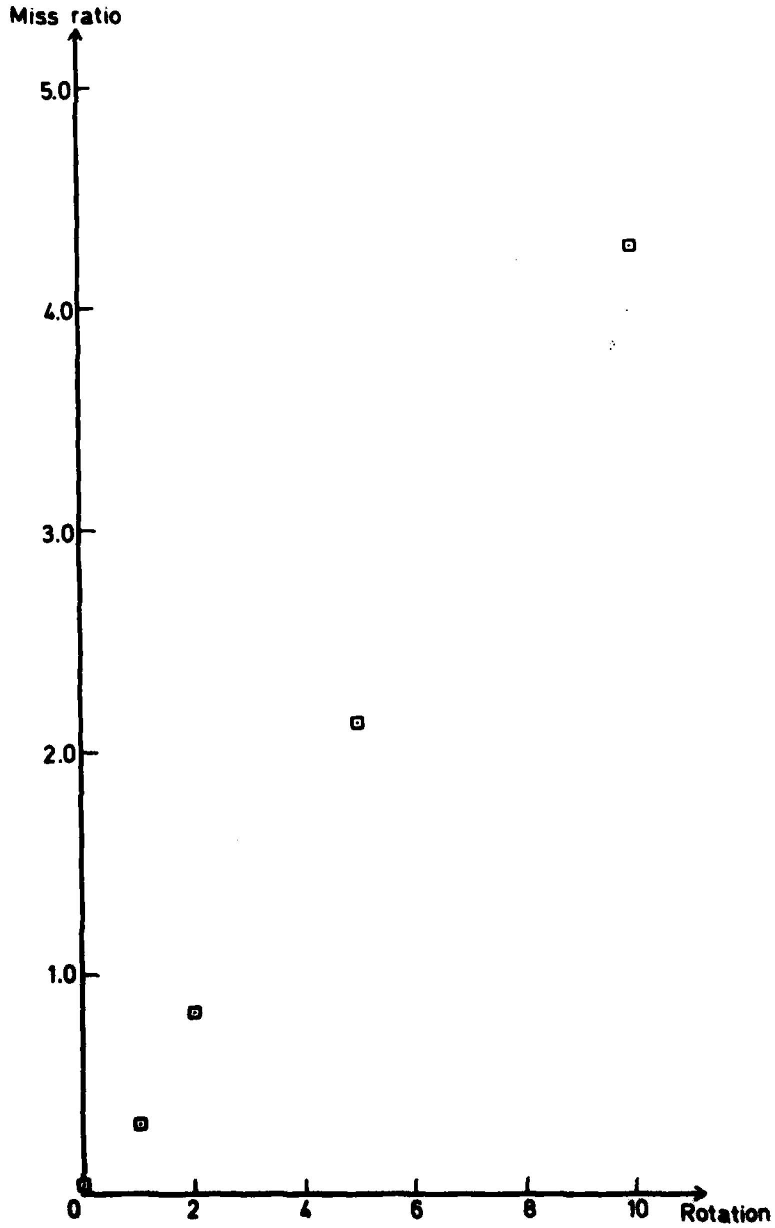

TR 79121

Fig 6 Miss ratio

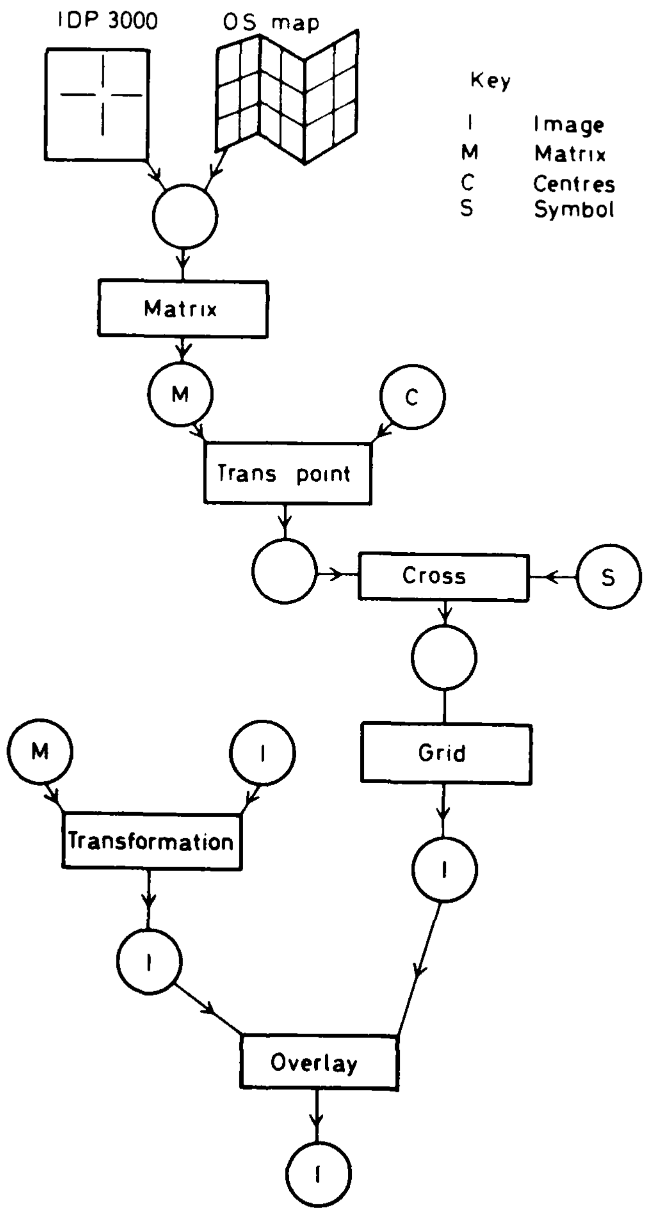

Fig 7 Software system diagram

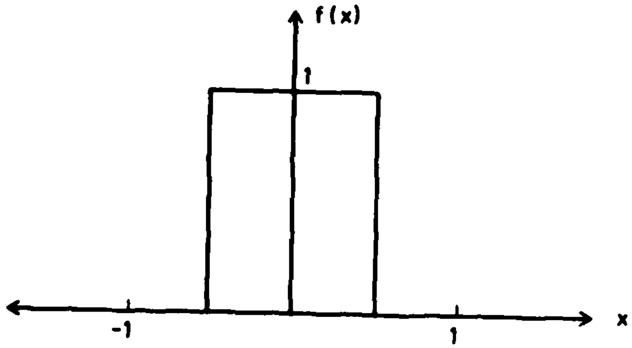

Fig 8a Nearest neighbour kernel

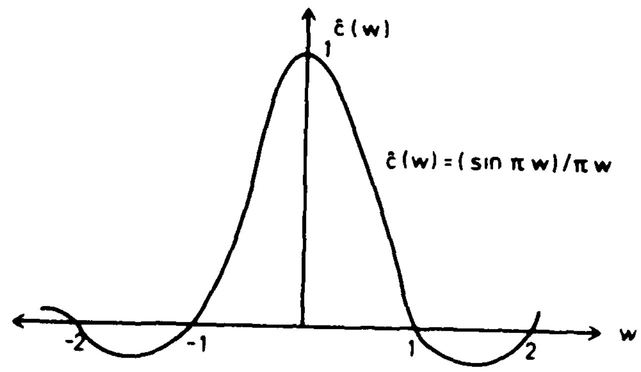

Fig 8b Nearest neighbour transform

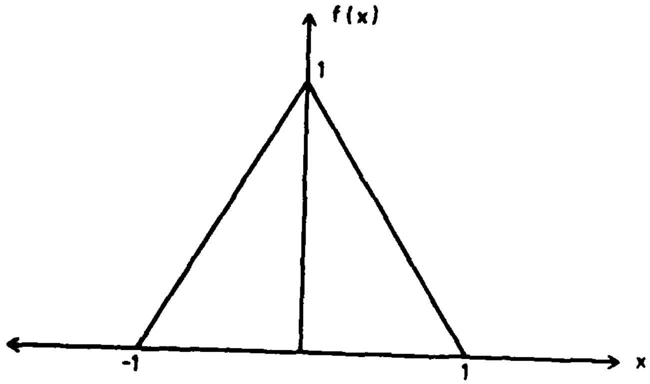

Fig 8c Linear kernel

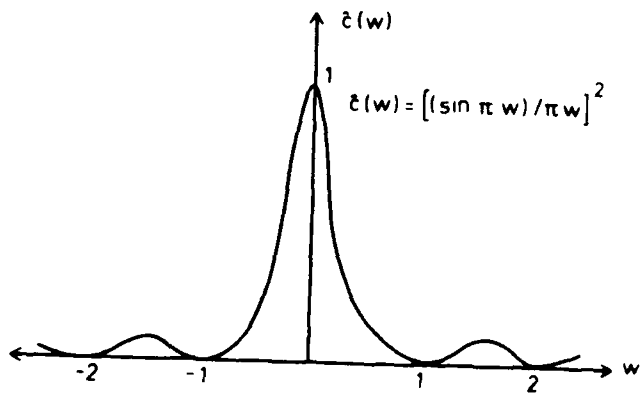

Fig 8d Linear transform

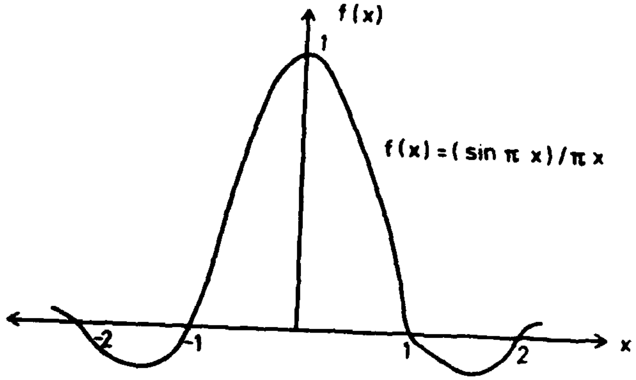

Fig 9a Ideal kernel

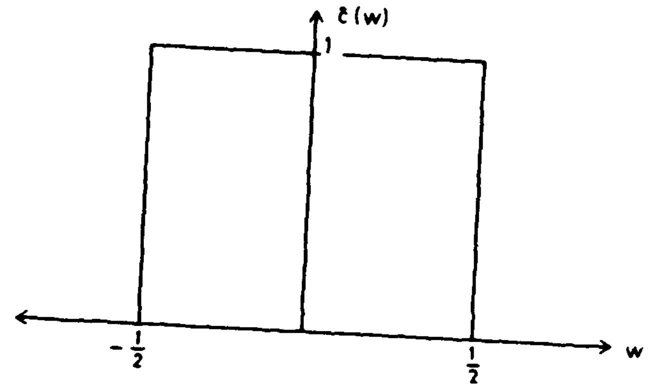

Fig 9b Ideal transform

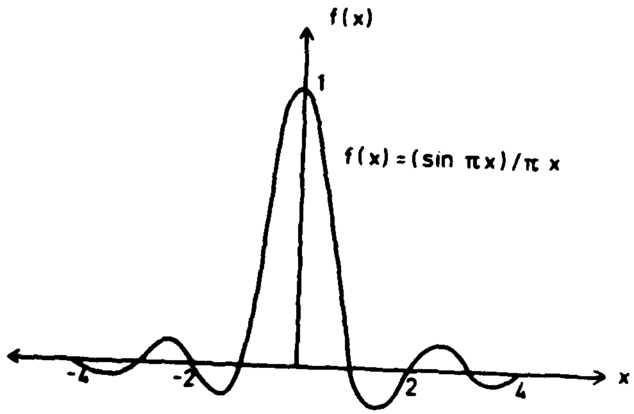

Fig 9c Eight-point truncation

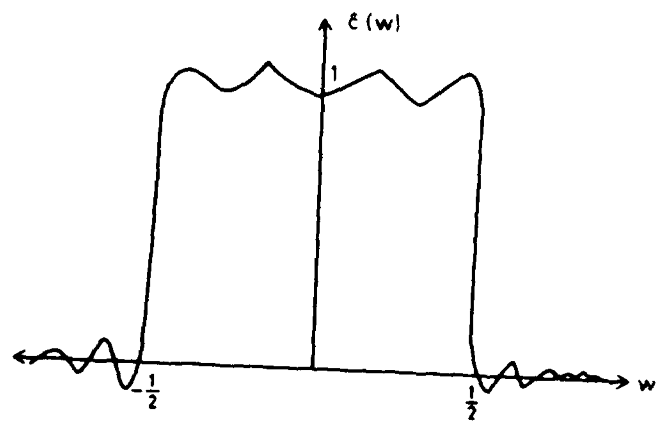

Fig 9d Eight-point transform

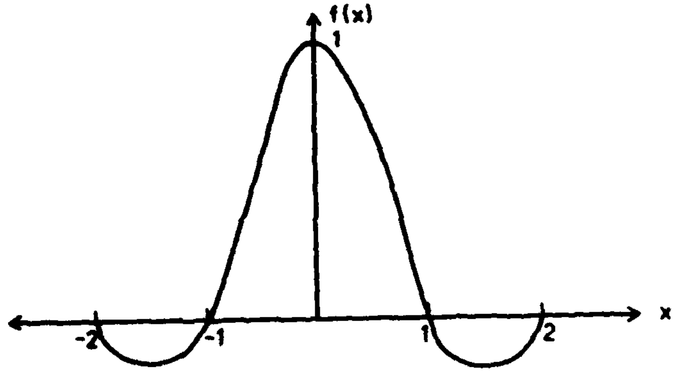

Fig 10a **Cubic kernel**

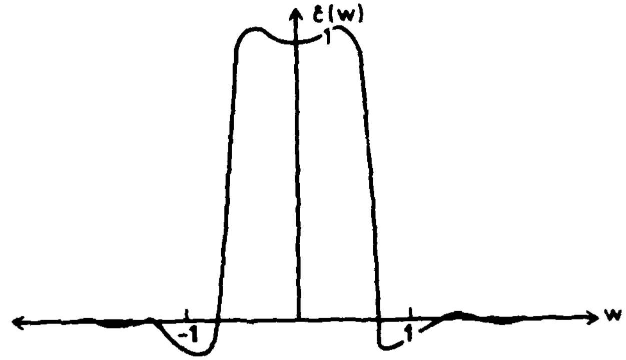

Fig 10b **Cubic transform**

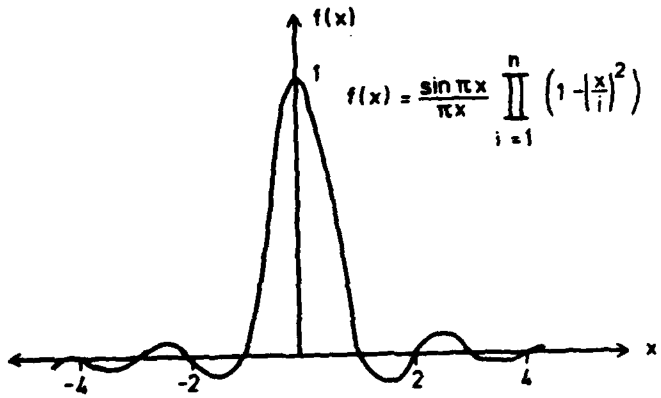

Fig 10c **Shlien kernel**

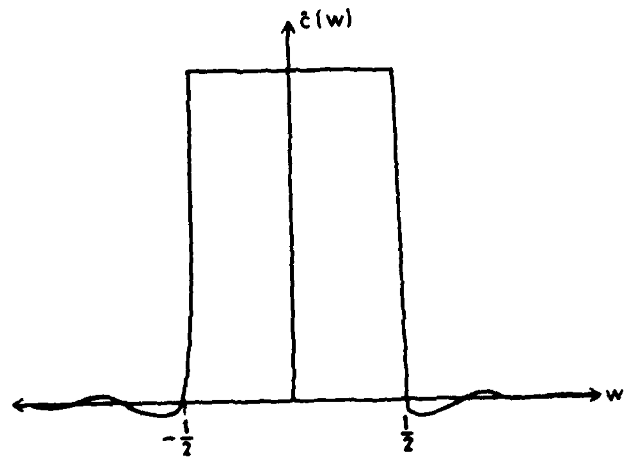

Fig 10d ***Shlien transform***

REPORT DOCUMENTATION PAGE

Overall security classification of this page

UNLIMITED

As far as possible this page should contain only unclassified information. If it is necessary to enter classified information, the box above must be marked to indicate the classification, e.g. Restricted, Confidential or Secret.

1. DRIC Reference (to be added by DRIC)	2. Originator's Reference	3. Agency Reference	4. Report Security Classification/Marking
	RAE TR 79121	N/A	UNLIMITED

5. DRIC Code for Originator	6. Originator (Corporate Author) Name and Location
7673000W	Royal Aircraft Establishment, Farnborough, Hants, UK
5a. Sponsoring Agency's Code	**6a. Sponsoring Agency (Contract Authority) Name and Location**
N/A	N/A

7. Title Geometric correction of satellite imagery

7a. (For Translations) Title in Foreign Language

7b. (For Conference Papers) Title, Place and Date of Conference

8. Author 1. Surname, Initials	9a. Author 2	9b. Authors 3, 4	10. Date	Pages	Refs.
Williams, J.M.			September 1979	30	4

11. Contract Number	12. Period	13. Project	14. Other Reference Nos.
N/A	N/A		Space 569

15. Distribution statement

(a) Controlled by – Head of Space Department, RAE (RAL)

(b) Special limitations (if any) –

16. Descriptors (Keywords) (Descriptors marked * are selected from TEST)

Remote sensing. Geometric correction. Landsat. Satellite imagery.

17. Abstract

[illegible]